Submitted to the Annals of Applied Statistics
arXiv: arXiv:0000.0000

arXiv:1403.5196v1 [stat.ME] 20 Mar 2014

CALIBRATION OF COMPLEX COMPUTER SIMULATORS USING LIKELIHOOD EMULATION

By Jeremy E. Oakley*, and Benjamin D. Youngman†

*University of Sheffield** *and University of Exeter*†

We calibrate a Natural History Model, which is a class of computer simulator used in the health industry, and here has been used to characterise bowel cancer incidence for the UK. To calibrate the simulator, we must address that it is stochastic and its output is count data. The simulator is of 'moderate' expense, which prohibits calibration by Monte Carlo sampling. To achieve a reliable calibration, we must also specify its discrepancy: how, when calibrated, it differs from reality. We propose a method for calibration that combines a statistical emulator for the likelihood function with importance sampling. The emulator provides an interim sample of inputs at which the simulator is run, from which the likelihood is calculated. Importance sampling is then used to re-weight the inputs and provide a final sample of calibrated inputs. Re-calculating the importance weights incurs little computational cost, and so we can easily investigate how different discrepancy specifications affect calibration.

1. Introduction. We consider the problem of calibrating a computer model of a physical process to observations from the process: finding model input values such that the model outputs match the observed data as closely as possible. Our approach is inspired by the framework for Bayesian calibration proposed by [11] and developed in [10], [3], [2] and [9], and by Bayes linear history matching developed in [5], [6] and [16]. In this paper, we refer to a computer model as a "simulator". We focus on three issues, which we discuss now in turn: computationally expensive simulators, "discrepancy": the error in a simulator prediction due to the simulator being an imperfect model of reality, and stochastic simulators: simulators that can return different output values when run repeatedly at the same input values.

Any calibration method will involve running the simulator at different input values, and so methods that require large numbers of simulator runs become impractical if a single simulator run at one input value takes a long time. A well-established technique for handling expensive simulators, proposed in [12], is to construct a cheap surrogate model or "emulator" of the simulator using Gaussian process regression, based on a relatively small

Keywords and phrases: Bayesian inference, calibration, computer simulator, statistical emulation, importance sampling

number of simulator runs. Variations of this method are used in the above references. In this paper, we consider a simulator of 'intermediate' computational cost, in which runs take between one and two minutes, depending on input values. We argue that this changes the nature of the surrogate modelling problem. In our proposed approach, rather than attempting to construct a very precise emulator of the simulator, we use a cruder emulator to guide us to the appropriate regions of the input space, and then do direct simulator evaluations in those regions. In particular, we propose the use of importance sampling, where the emulator is used to construct the importance density.

When calibrating a simulator, it is important to account for simulator discrepancy for two reasons. Firstly, if the inputs are physically meaningful quantities that could, in principle, be observed directly, calibrating a simulator without accounting for discrepancy may result in biased estimates with severe over-confidence, as demonstrated in [4]. If the simulator inputs are 'tuning' parameters that are not physically observable, discrepancy plays an important role when calibrating to multiple outputs, or when we wish to predict unobserved output quantities using a calibrated simulator. Suppose that we have a physical observation for an output quantity Z_1, and wish to predict an unobserved output quantity Z_2. A simulator input value may give a poor fit to output Z_1, but a good prediction of Z_2. If we do not believe the simulator models Z_1 perfectly, we would not necessarily want to rule out such an input value and corresponding prediction of Z_2. Accounting for the simulator error in modelling Z_1 would prevent this.

As argued in [4], it is important to specify meaningful proper prior distributions for simulator discrepancy, but to do this may be difficult. In [17], within a Bayes linear framework, the simulator expert only provided an *interval* for the variance of a discrepancy parameter. [13] suggest 'opening the black box' and incorporating discrepancy terms within the simulator, so that the expert considers sources of simulator discrepancy explicitly, rather than attempting to make judgements about the overall discrepancy. We argue that it is desirable to be able to investigate, without too much difficulty, a range of different discrepancy distributions, within any calibration methodology. Within our proposed importance sampling framework, we suggest an initial, conservative specification of simulator discrepancy, which can then be varied with little extra computational effort via re-calculation of importance weights corresponding to different discrepancy distributions.

The final issue we consider is that of a stochastic simulator, which raises the question of what it is we should be trying to emulate, assuming that an emulator is necessary. The simulator in our case study produces random

count data. In a similar scenario, [8] constructed emulators for probabilities from which the count data were assumed to have been generated. Here, we propose constructing an emulator for the likelihood function given the observed data. Our simulator produces 30 count data outputs (with various dependencies between the outputs), and so emulating the likelihood reduces the computational effort to emulating a univariate output, and enables us to implement an importance sampling approach for the calibration.

This paper has the following structure. The next section outlines the calibration method. Section 3 presents the results of calibration of a Natural History Model and section 4 offers conclusions and discussion of the calibration method.

2. Outline method.

2.1. *The calibration problem.* We have target data Z, observed in the real world, with which we can calibrate the simulator. The data are made up of observations of various binomial and multinomial random variables, but to simplify the discussion, we suppose that Z is a single binomial random variable, with $Z|\theta^* \sim Bin(N, \theta^*)$. The computer simulator encodes a function $\theta(x)$ that describes the relationship between some input parameters x and a binomial distribution probability parameter $\theta(x)$. We suppose that there is a true, observable input value X, observable in the sense that, in theory data could be obtained to estimate X directly, independently of the simulator. Relating the simulator to reality, we recognise that the simulator is not perfect, so that $\theta^* = \theta(X) + \delta$, where δ represents the simulator error or discrepancy. The calibration problem is to infer X given Z.

2.2. *Calibrating a stochastic computer simulator.* The computer simulator does not actually output $\theta(x)$ for a given input x. Instead, the simulator outputs a random variable $Y(x)$ with $Y(x)|\theta(x), n(x) \sim Bin(n(x), \theta(x))$. The value of $n(x)$ is expected to increase with the patient cohort size, the original patient sample size chosen for the simulator, but is subject to some random variation. Hence, for any simulator run at input x, we will have to infer the value of $\theta(x)$ based on the observations for $Y(x)$ and $n(x)$. During the calibration process, we will run the simulator at inputs $x_1, \ldots, x_m$, to obtain simulator data $D = \{x_i, Y(x_i), n(x_i)\}_{i=1}^m$, and so the aim of the calibration is to derive the posterior distribution $\pi(X|Z, D)$; we infer X given Z *and* D.

We can evaluate the likelihood $\pi(Z|x, D)$ for X at the value of x via

$$\pi(Z|x, D) = \iint \pi(Z|\theta(x), \delta, x, D)\pi(\theta(x)|x, D)\pi(\delta|\theta(x), x, D)d\theta(x)d\delta,$$

which we assume can be simplified as

$$\pi(Z|x,D) = \iint \pi(Z|\theta(x),\delta)\pi(\theta(x)|D)\pi(\delta|\theta(x))d\theta(x)d\delta.$$

We make a further simplification: we suppose that we have run the simulator at x to observe $Y(x)$ and $n(x)$, so that $\{x, Y(x), n(x)\} \in D$ and then we set

$$\pi(\theta(x)|D) = \pi(\theta(x)|Y(x), n(x)),$$

so that we only use the run at x to infer the corresponding $\theta(x)$.

2.3. *Incorporating simulator discrepancy.* As we have already discussed, specifying a single choice of discrepancy distribution is difficult, and so we propose the following strategy to account for simulator discrepancy. We start with a conservative prior distribution for δ that permits moderately large values. We obtain a sample from the posterior distribution $\pi(X|Z,D)$. We can then explore alternative distributions for δ, using importance sampling to re-weight the sample according to alternative prior distributions $\pi(\delta)$. For example, in the case where δ is multivariate, corresponding to a multiple output simulator, we can investigate scenarios where some outputs are believed to be better modelled than others. By starting with a conservative prior for δ we are, in effect, 'broadening the search' for inputs that give simulator outputs that are close to the observed data. Without any discrepancy, it is possible that no input value will give a good fit to all the output data.

Writing $\theta^* = \theta(X)+\delta$, we want the discrepancy term δ to add uncertainty about θ^* given $\theta(X)$, as we don't believe that running the simulator at the true observed X (and an infinitely large cohort of patients) would give us θ^*. To simplify the computation, we can achieve this effect by inflating the uncertainty about $\theta(x)$ given $n(x)$ and $Y(x)$, rather than by introducing an additional term δ. We choose a $U[0,1]$ prior distribution for $\theta(x)$ and suppose that

$$\theta(x)|Y(x), n(x) \sim Beta(1+\lambda Y(x), 1+\lambda(n(x)-Y(x))),$$

with $\lambda \in (0,1)$. The parameter λ has the effect of allowing for simulator discrepancy, by downweighting the information that the simulator run gives us about $\theta(x)$. In section 3.8 we investigate the sensitivity of calibration to different choices of λ. Using the distribution for $\theta(x)$ given in section 2.3, we assume $\delta = 0$ and re-write the likelihood as

$$\begin{aligned}\pi(Z|x,D) &= \int \pi(Z|\theta(x))\pi(\theta(x)|Y(x),n(x))d\theta(x) \\ &= \frac{{}^{N}C_{Z}B(1+\lambda Y(x)+Z, 1+\lambda(n(x)-Y(x))+N-Z)}{B(1+\lambda Y(x), 1+\lambda(n(x)-Y(x)))},\end{aligned}$$

where $B(.,.)$ is the Beta function.

2.4. *Sampling from the posterior distribution of the inputs.* Obtaining $Y(x)$ and $n(x)$ is computationally expensive, so we need to be selective in when we choose to run the simulator and evaluate the likelihood. We use importance sampling, where we construct a cheap-to-evaluate importance density using a Gaussian process emulator [12]. This will be an iterative procedure, as the initial design region for the inputs may be specified somewhat conservatively, so that it may take several attempts to construct a satisfactory importance density. The outline procedure is as follows.

1. Using an initial set of simulator runs, investigate the design region to see if any subregions can be ruled out as having relatively low likelihood.
2. Run the simulator at a moderate number of input values over the reduced design region, to get initial simulator data $D = \{x_i, Y(x_i), n(x_i)\}_{i=1}^{m}$. Evaluate the likelihood in equation (2.1) for each input value x_i.
3. Using the data from Step 2, construct a fast approximation of the log-likelihood using a Gaussian process emulator.
4. Construct an importance density for $\pi(X|Z)$ by approximating the log-likelihood by the posterior mean of the emulator. Use MCMC to generate a sample of values $X_1, \ldots, X_r$ from this approximate posterior density. To guard against the support of the importance density being too small, flatten the log-likelihood by multiplying it by a suitable constant.
5. Run the simulator at $X_1, \ldots, X_r$ and evaluate the likelihood (2.1) for each of these points. Calculate importance weights for each input.
6. If a small proportion of the inputs in $X_1, \ldots, X_r$ have relatively large importance weight, update the emulator to include the likelihood evaluations, and return to Step 4.

3. Calibration of a Natural History Model.

3.1. *Natural History Models.* Here we calibrate a stochastic computer simulator called a Natural History Model (NHM), which was developed in [14] to characterise UK bowel cancer incidence. It has 25 calibration parameters, X, which are unknown; some, however, are in principal physically observable. We will calibrate the NHM to derive the joint distribution of X given target data. For brevity we only describe the essence of the NHM's workings here; see [14] for a fuller description.

The NHM represents a *birth cohort*: a fixed-size sample of people followed from birth to death. A person in the cohort is deemed to have devel-

oped bowel cancer when they have reached the first cancer state, Duke's A, having begun in a non-cancer state, and progressed through three, ordered pre-cancer states: low-, medium- and high-risk adenomas. A person may continue to progress through three more increasingly severe cancer states, Duke's B, C and Stage D. Progression between states is governed by time. When in a given state, a progression time to the next state is simulated, together with a presentation time (the most common form of presentation being to visit a doctor), and a time until death. Out of these three actions, the one that occurs is the one with the shortest simulated time. Times follow state-dependent Weibull distributions, the parameters of which form the majority of the NHM's unknown parameters that we calibrate. This framework allows a person's age to be known whenever they change state. It also allows a person to progress straight from birth to death (without ever contracting bowel cancer), or to progress through some or all pre-cancer and cancer states. By presenting a patient enters the health system where they receive a bowel cancer diagnosis. The age-based data that form part of the NHM's output result from these diagnoses and the tracking of ages. Having left the health system, a person is still represented by the NHM, but their progression rates between states are elevated. The following gives details of the NHM's output required for calibration.

3.2. *Target data, output and notation.* The target data and NHM output are counts that we will in general denote by Z_{jk} and $Y_{jk}(x)$, respectively, where $j = 1, \ldots, 4$ indexes the data type and $k = 1, \ldots, K_j$ indexes groups within types; corresponding sample sizes are denoted N_{jk} and $n_{jk}(x)$, respectively. Here x is the input vector that we use to initialise the NHM. The data types are identified explicitly, as opposed to considering the output as a single vector, due to their inherent differences, which will emerge in the following summaries.

3.2.1. *Cases by age.* Target data Z_{1k} represent a cross-sectional study and give the number of people out of N_{1k} in the UK developing bowel cancer in 2008, where $k = 1, \ldots, 18$ indexes age groups 0-4, 5-9, ..., 80-84, 85+. (The data are freely available[1].) The NHM's output does not match the target data directly. Instead, it represents a birth cohort comprising P people, by modelling the cancer state and age of each person. To make the NHM output consistent with the target data, it is resampled by allocating each person to age group $k = 1, \ldots, 18$ at random, according to probabilities

[1]Data can be downloaded from http://info.cancerresearchuk.org/cancerstats/types/bowel/incidence/

determined by proportions in the UK population. Let $r = 1, \ldots, R$ index each randomisation. The resulting NHM output corresponding to Z_{1k} is denoted $Y_{1k}^{(r)}(x)$, with corresponding sample size $n_{1k}^{(r)}(x)$. The likelihood is obtained by averaging over randomisations, with R large.

3.2.2. *Cases by type.* Z_{2k} is the number of bowel cancer cases of type k out of N_2 cases, where $k = 1, \ldots, 4$ indexes types Duke's A, B and C, and Stage D, respectively. The NHM output is denoted $Y_{2k}(x)$ and is directly comparable to Z_{2k}. The total number of cases simulated is denoted $n_2(x)$.

3.2.3. *Obstructed cases by type.* These data also represent cases by type, but only those cases in which an obstruction occurs and only for types Duke's B, C and Stage D [15]. We therefore define Z_{3k}, N_3, $Y_{3k}(x)$ and $n_3(x)$ similarly to $j = 2$.

3.2.4. *Undetected adenomas by age.* Z_{4k} is the number of people out of N_{4k} that have developed adenomas, but in whom they have not been detected, where $k = 1, \ldots, 4$ indexes age groups under 55, 55-64, 64-74 and over 75 [18]. NHM output $Y_{4k}(x)$ and $n_{4k}(x)$ are defined similarly.

3.3. *Discrepancy specification.* To introduce simulator discrepancy to the NHM, we consider reducing output sample sizes and counts, $n_{jk}(x)$ and $Y_{jk}(x)$, and specify these reductions as fractions, $\lambda_j \in (0, 1)$, $j = 1, \ldots, 4$. To assess calibrated output, we consider its similarity to the target data, given approximate error bounds. These bounds represent how close a simulator output should be to the target data, considering three sources of error: sampling variability in the data, stochastic variability of the simulator output, and simulator discrepancy. For brevity, we present results for binomial data, though only minor alterations are required for multinomial data.

We first consider error due to sampling variability. If $Z|\theta^* \sim Bin(N, \theta^*)$, then the variance of $p := Z/N$, which is used to estimate θ^*, is $p(1-p)/N$. Similarly, if $Y(x) \sim Bin(n(x), \theta(x))$ is simulator output without discrepancy, the estimator $p(x) := Y(x)/n(x)$ has variance $p(x)(1-p(x))/n(x)$. The addition of simulator discrepancy, through $\lambda \in (0, 1)$, inflates the variance of the estimator to $p(x)(1-p(x))/(\lambda n(x))$, which can be partitioned as

$$\frac{p(x)(1-p(x))}{\lambda n(x)} = \frac{p(x)(1-p(x))}{n(x)} + \frac{p(x)(1-p(x))(1-\lambda)}{\lambda n(x)}.$$

Thus we decompose the variance of the simulator output into contributions due to the simulator being stochastic and that added by it being imperfect.

We assess the calibrated output against the target data by considering approximate 95% intervals around the target data, which widen as we add in the different sources of error:

$$\begin{array}{ll} \text{measurement error} & \pm 2\sqrt{\dfrac{p(1-p)}{N}}, \\ \text{measurement error and simulator uncertainty} & \pm 2\sqrt{\dfrac{p(1-p)}{N} + \dfrac{p(x)(1-p(x))}{n(x)}}, \\ \text{measurement error, simulator uncertainty and simulator discrepancy} & \pm 2\sqrt{\dfrac{p(1-p)}{N} + \dfrac{p(x)(1-p(x))}{n(x)} + \dfrac{p(x)(1-p(x))(1-\lambda)}{\lambda n(x)}}. \end{array}$$

While $p(x)$, $n(x)$ and $Y(x)$ vary with x, they are estimated only once, from the simulator run with highest likelihood.

Figure 1 displays variance decompositions for each data source. This visual representation allows us to choose values of λ_j "by eye": we choose values to give bounds around the target data that are such that, if output falls within the bounds, then we are prepared to deem it and its corresponding input plausible. We perform the calibration in waves and, in the build-up to the final calibration, can broaden the search for inputs by extending these intervals. We investigate sensitivity to different choices of λ in section 3.8. In particular, our method is intended to make such sensitivity analyses relatively straightforward. Ultimately we set $\lambda_1 = 0.8$, $\lambda_2 = 0.008$, $\lambda_3 = 0.04$ and $\lambda_4 = 0.0004$, which are the values represented in Figure 1.

3.4. *Likelihoods for the cancer data.* Combining sections 2 and 3.2 allows us to calculate the likelihood for all the NHM's output. Notation for realisations follows from section 3.2; for example, z_{1k} is the observed number of people in age group k developing bowel cancer out of N_{1k} and $y_{1k}(x_i)$ is the corresponding NHM count out of $n_{1k}(x_i)$ for input x_i, with notation for other data types defined similarly. We model the cases by age and undetected adenomas by age data as binomially distributed, and assume weak prior information for its parameters by adopting a Uniform[0,1] prior. We assume the cases by type and obstructed cases by type data to be multinomially distributed, and use a Dirichlet($\mathbf{1}$) prior to again represent weak prior knowledge. Finally the complete target data are $z = (z_1, z_2, z_3, z_4)$ where $z_j = (z_{j1}, \ldots, z_{jK_j})$.

The overall log-likelihood for the complete target data for an input x_i at

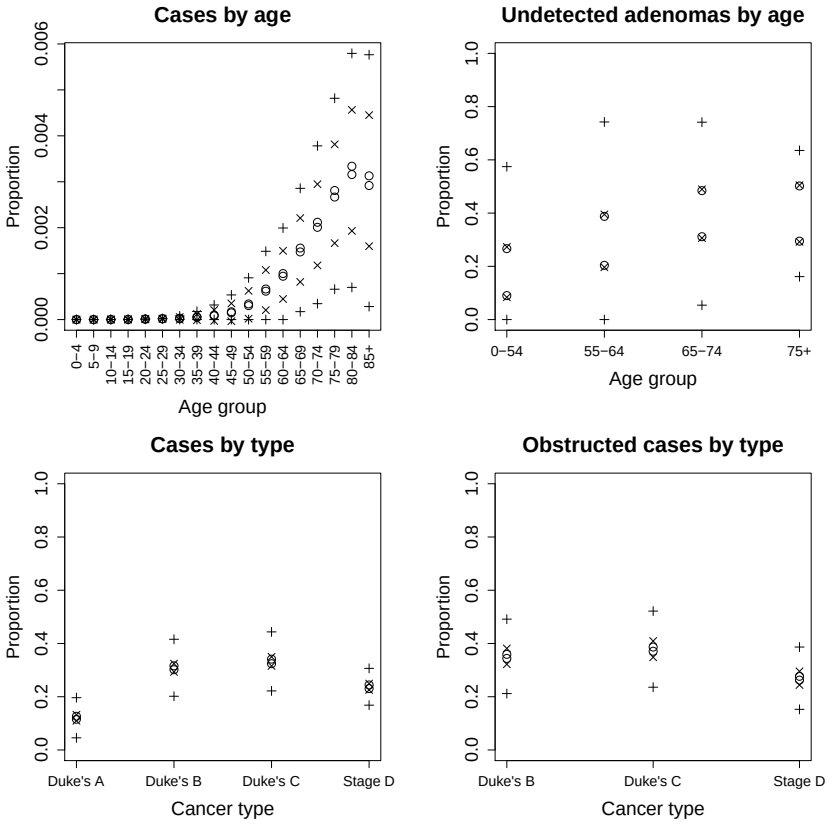

FIG 1. *Variance decompositions for each target data source described in section 3.2. Contributions to variability due to target data (∘), simulator uncertainty (×) and simulator discrepancy (+) are shown.*

which we have run the simulator and obtained output $y(x_i)$ is given by

$$\log\{\pi(z \mid x_i,\, y(x_i))\} = \sum_{j=1}^{4} \log(\pi_j), \tag{3.1}$$

where

$$\pi_1 = \frac{1}{R}\sum_{r=1}^{R}\left\{\prod_{k=1}^{K_1} \frac{N_{1k}!\, B(1+z_{1k}+\lambda_{1k}y_{1k}^{(r)}(x_i),\, 1+N_{1k}-z_{1k}+\lambda_{1k}\{n_{1k}^{(r)}(x_i)-y_{1k}^{(r)}(x_i)\})}{(N_{1k}-z_{1k})!\, z_{1k}!\, B(1+\lambda_{1k}y_{1k}^{(r)}(x_i),\, 1+\lambda_{1k}\{n_{1k}^{(r)}(x_i)-y_{1k}^{(r)}(x_i)\})}\right\}$$

with index $r = 1, \ldots, R$ denoting the rth randomisation of the NHM output,

$$\pi_2 = \frac{N_2!\,\{\lambda_2 n_2(x_i) + K_2 - 1\}!}{\{N_2 + \lambda_2 n_2(x_i) + K_2 - 1\}!} \prod_{k=1}^{K_2} \frac{\{z_{2k} + \lambda_2 y_{2k}(x_i)\}!}{z_{2k}!\{\lambda_2 y_{2k}(x_i)\}!},$$

$$\pi_3 = \frac{N_3!\,\{\lambda_3 n_3(x_i) + K_3 - 1\}!}{\{N_3 + \lambda_3 n_3(x_i) + K_3 - 1\}!} \prod_{k=1}^{K_3} \frac{\{z_{3k} + \lambda_3 y_{3k}(x_i)\}!}{z_{3k}!\{\lambda_3 y_{3k}(x_i)\}!},$$

$$\pi_4 = \left\{ \prod_{k=1}^{K_4} \frac{N_{4k}!\,B(1 + z_{4k} + \lambda_{4k} y_{4k}(x_i),\ 1 + N_{4k} - z_{4k} + \lambda_{4k}\{n_{4k}(x_i) - y_{4k}(x_i)\})}{(N_{4k} - z_{4k})!\,z_{4k}!\,B(1 + \lambda_{4k} y_{4k}(x_i),\ 1 + \lambda_{4k}\{n_{4k}(x_i) - y_{4k}(x_i)\})} \right\}.$$

We calculate the log-likelihood for 10,000 NHM runs, each using a birth cohort of size 100,000. Figure 2 shows the log-likelihood against inputs 1, 2, 3, 12, and 25, specifically against single inputs (achieved by maximising the likelihood over equal-sized bins) and for pairwise combinations of inputs (achieved by maximising over grid cells). Input 1 represents the age at which a person can develop adenomas, input 2 the log-parameterised Weibull shape parameter for progression times between pre-cancer states, input 3 the Weibull scale parameter for progression to the first pre-cancer state, input 12 the change in Weibull scale parameters due to having previously been treated for cancer and input 25 the probability of a person developing cancer.

Figure 2 shows that for some regions of input space the log-likelihood is much higher than for others. We use where the likelihood is relatively high to define a reduced input space, which is specified by marginal ranges and pairwise regions. Because we start with broad parameter ranges for all 25 inputs, there is large variation in the likelihood values of Figure 2, and so our criterion for ruling out parts of input space is set conservatively: we omit parts where the likelihood ratio, relative to the observed maximum, fails to exceed e^{-40}. This reduces the input space to 0.7% of its original size. As we approximate true maximum log-likelihoods by those observed, we make conservative choices here to compensate for observed maxima being underestimates of the true maxima. This could be avoided if it were feasible to use many more simulator runs. This technique for reducing the input region is related to that used in history matching by [16], in which implausibility of parts of input space is quantified, and parts measured to have large implausibility are ruled out.

As in [16], the input region can be further reduced in waves. Here second and third waves, also of 10,000 NHM runs, are performed, which use birth cohorts of 200,000 and 300,000 people, respectively. The emulator training data are based on the region chosen after the third wave, which is 0.0001%

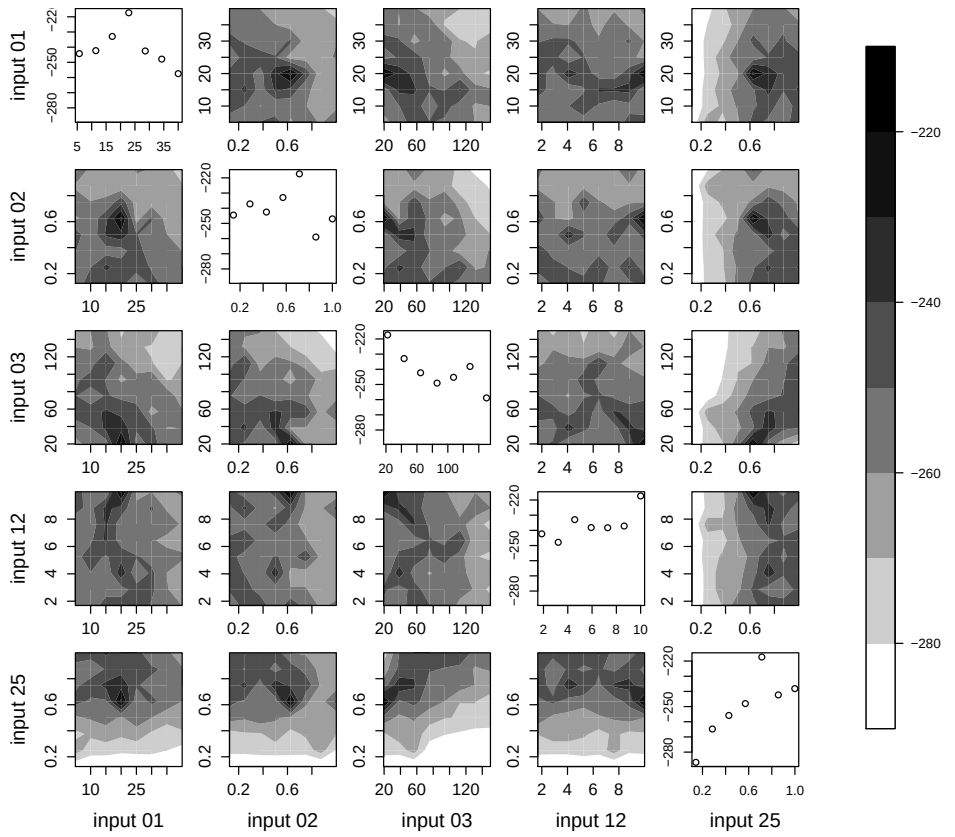

FIG 2. *Pairwise maximised log-likelihood (off-diagonal) and marginal binned maximised log-likelihoods (diagonal) for inputs 1, 2, 3, 12 and 25.*

the size of our starting input region.

3.5. *Emulator specification and building.* We are building an emulator for the function $f(x)$, the log-likelihood for input vector x, where $x = (x^{(1)}, \ldots, x^{(p)})^T$, which is defined in section 3.4 as

$$f(x) := \log\{\pi(z \,|\, x,\, y(x))\} = \sum_{j=1}^{4} \log \pi_j.$$

Thus we model

$$f(x) \,|\, \sigma^2, \beta, \phi, \nu^2 \sim GP(h^T(x)\beta, \sigma^2 c(x,\,))$$

where $h(\,)$ and β comprise q basis functions and regression coefficients, respectively, $h^T(x)\beta$ is therefore the GP mean function, σ^2 is its variance and

$c(\,,\,)$ is its correlation function. We choose the correlation function to have the Gaussian form

$$c(x_i, x_j) = \begin{cases} \exp\{-\sum_{d=1}^{p}((x_i^{(d)} - x_j^{(d)})/\phi_d)^2\} & \text{if } x_i \neq x_j, \\ (1 + \nu^2/\sigma^2)^{-1} & \text{if } x_i = x_j, \end{cases}$$

for a set of roughness parameters $\phi = \{\phi_1, \ldots, \phi_p\}$, where $\phi_d > 0$, $d = 1, \ldots, p$. The parameter $\nu^2 > 0$ introduces a *nugget* effect into the emulator, which has been shown to improve the predictive performance of Gaussian process emulators [1, 7], but is imperative for a stochastic simulator such as the NHM. The emulator is specified to have a constant mean function, ie. $h(x) = 1$. Polynomial terms could be added, but perhaps more suitable would be (the log of) a parametric density function, though this gives a mean function that is non-linear in its parameters.

We use 2,000 simulator runs for the emulator training data, which are chosen using a Maximin Latin hypercube design on the reduced region chosen after wave three in section 3.4. We define the following: input set $D_X = \{x_1, \ldots, x_m\}$; vector of corresponding log-likelihoods $f(D_X) = (f(x_1), \ldots, f(x_m))^T$; $m \times m$ matrix A, which has (i, j)th element $c(x_i, x_j)$; and $t(x)^T = (c(x_1, x), \ldots, c(x_m, x))$.

For the hyperparameter prior we choose $\pi(\sigma^2, \beta, \phi, \nu^2) \propto \sigma^{-2}$. It follows that posterior emulator is given by

$$f(x) \,|\, D, \phi, \nu^2 \sim tP_{n-q}(\hat{\beta}, \hat{\sigma}^2 c^*(x, \,)),$$

a Student t-process on $n - q$ degrees of freedom, where

$$\begin{aligned} \hat{\beta} &= (1_m^T A^{-1} 1_m)^{-1} 1_m^T A^{-1} f(D_X) \\ \hat{\sigma}^2 &= (m - q - 2)^{-1} (f(D_X) - \hat{\beta})^T A^{-1} (f(D_X) - \hat{\beta}) \\ m^*(x) &= \hat{\beta} + t(x)^T A^{-1} (f(D_X) - \hat{\beta}) \end{aligned}$$

and

$$\begin{aligned} c^*(x, x') = c(x, x') &- t(x)^T A^{-1} t(x') \\ &+ (1 - t(x)^T A^{-1} 1_m)(1_m^T A^{-1} 1_m)^{-1} (1 - t(x')^T A^{-1} 1_m)^T. \end{aligned}$$

Finally, (ϕ, ν^2) has posterior

$$\pi^*(\phi, \nu^2) \propto (\hat{\sigma}^2)^{-(m-q)/2} |A|^{-1/2} |1_m^T A^{-1} 1_m|^{-1/2} \, \pi(\phi, \nu^2).$$

We fix (ϕ, ν^2) at the mode of $\pi^*(\phi, \nu^2)$. This is found using the Nelder-Mead optimisation algorithm, which is initialised with 200 iterations of the Gibbs sampler, in which Metropolis-Hastings updates are used.

3.6. *Input sampling.* For the algorithm of section 2.4 to perform well, the emulator should represent high values of the log-likelihood fairly accurately. We use importance sampling to give a sample of inputs, and for the importance density use the emulator posterior mean, which serves as an approximation to the log-likelihood. We can sample from the importance density by again using Gibbs sampling with Metropolis-Hastings updates. The following algorithm then describes how we obtain the final sample of calibrated inputs. The algorithm is designed to identify parts of the input region where the difference between the posterior mean and the log-likelihood is large, or where, given the posterior mean is relatively big, the emulator's uncertainty is large. We can then add simulator runs in these parts to enable the emulator to provide a more accurate representation of the log-likelihood surface.

1. Obtain a sample of inputs, $D_S = (X_1, \ldots, X_S)$, by Gibbs sampling using the emulator posterior mean, $m^*(x)$, to approximate the log-likelihood.
2. Compute the pivoted Cholesky decomposition of the covariance matrix for the sample, ie. the $S \times S$ matrix A_S with (i,j)th element $c(X_i, X_i)$, $i, j = 1, \ldots, S$, and let $\{p_s\}_{s=1}^S$ denote its diagonal elements. Sort D_S by the pivot, and take the first u members, to give D_{piv}, where u is the maximum number of simulator runs we are prepared to add to the training data in one iteration.
3. Define p_s to be 'large' if $p_s > v$, for some $v > 0$. If no p_s are large, proceed to Step 5. Otherwise form the set $D^\dagger = \{X_s \in D_{piv} : p_s > 2\}$, for $s = 1, \ldots, u$, evaluate the simulator at each of its members and calculate their log-likelihoods, $f(D^\dagger)$.
4. Add $D^\dagger$ and $f(D^\dagger)$ to the training data, re-build the emulator, and return to Step 1.
5. Compute importance weights $w_s = \exp\{f(X_s) - m^*(X_s)\}$ for $X_s \in D_S$. If a large proportion of weights are zero, return to Step 4.
6. Obtain the calibrated inputs, $D^* = \{X_1^*, \ldots, X_M^*\}$, by resampling D_S with replacement according to weights $w_s^* = w_s / \sum_{s=1}^u w_s$.

For Step 1 of the calibration algorithm we choose $S = 2,000$, which is achieved by thinning an initial sample of size 100,000 by 50. For Step 2 we choose $u = 200$ and for Step 3 choose $v = 2$. During the first iteration of the algorithm we find that almost all p_s are large, which indicates that the emulator's uncertainty is large for all the sampled inputs. Consequently, the importance density may have insufficient support where the true log-likelihood is high. We flatten the log-likelihood to compensate for this, which

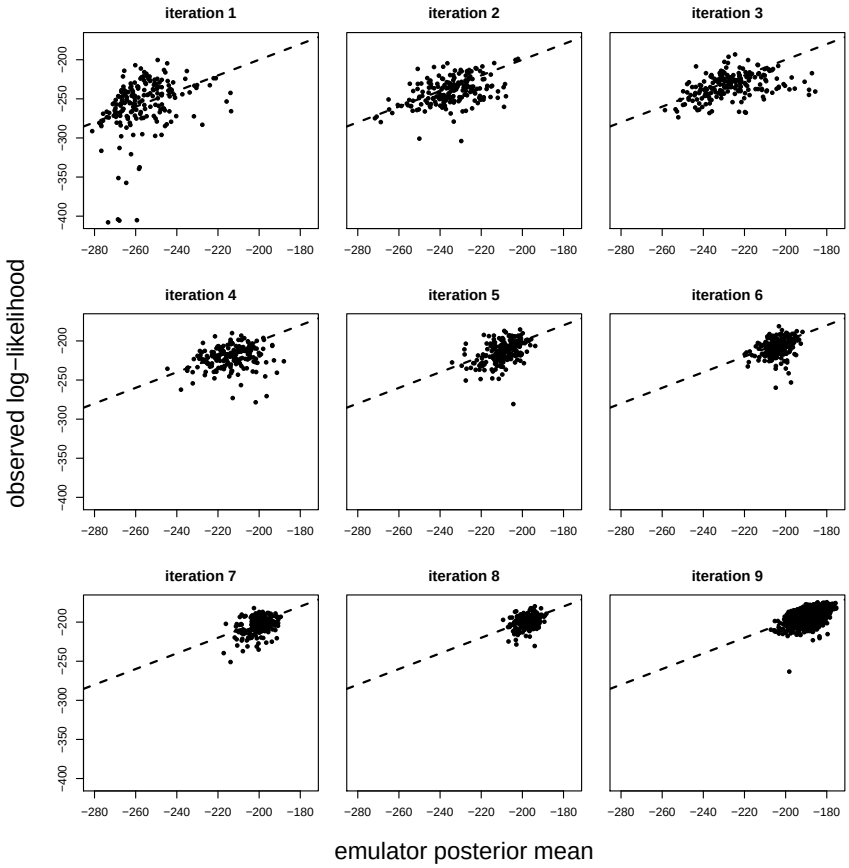

FIG 3. *Observed log-likelihoods against emulator posterior means (based on previous iteration) at iterations 1–8 for samples of size 200 and iterations 9 for a sample of size 1000. The line $y = x$ is superimposed (- - -).*

is achieved by using $\alpha m^*(x)$ instead of $m^*(x)$, $0 < \alpha \leq 1$, in Step 1; we initially choose $\alpha = 0.1$. Log-likelihoods calculated for the simulator runs are then compared against previous emulator posterior means, that is comparing $f(x)$ with $E(f(x) \mid D)$ for $x \in D^\dagger$, where D are the last-used training data. This comparison is shown for iterations 1–9 in Figure 3.

From Figure 3, we see that the agreement between $f(x)$ and $E(f(x) \mid D)$ is poor for the first iteration, which means that the emulator posterior mean will not serve well as an importance density for sampling inputs from the log-likelihood. We also look at how the simulator's output compares with the target data, given expected levels of uncertainty (as described in section

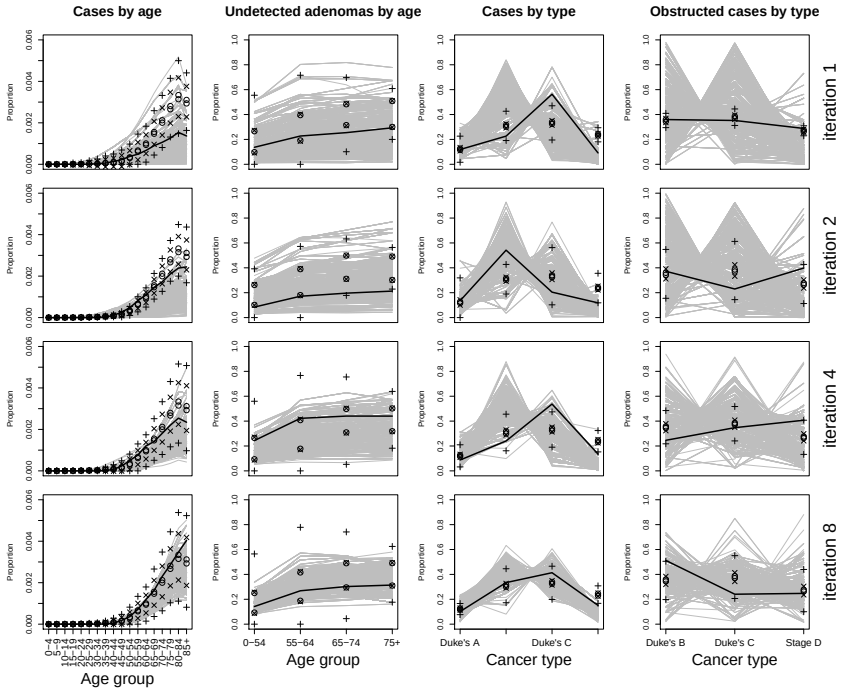

FIG 4. *NHM output against target data for iterations 1, 2, 4 and 8. Uncertainty bounds are as in Figure 1. The black line highlights the run with highest likelihood.*

3.3), which is shown for iterations 1, 2, 4, and 8 in Figure 4. For iteration 1, while some runs give a good match to some of the target data, most fail to provide an adequate match to all of the target data.

We proceed to perform further iterations. For iteration 2 we increase α to 0.2, and find that the match between $f(x)$ and $E(f(x) \mid D)$ has improved, but is still unsatisfactory, which can be seen in Figure 4. Therefore we perform further iterations, increasing α by 0.1 for each. Adequate agreement between the emulator and observed log-likelihoods is achieved by iteration 8, which is confirmed by iteration 9, the latter of which we choose to be the final emulator. There is some suggestion from Figure 4 of disagreement between the NHM output and the target data at iteration 8; however, the points used to assess this are those for which the emulator's conditional variance is greatest, and therefore a better match between the emulator's posterior mean and the true log-likelihoods can be expected for a random sample of inputs. Furthermore, we only need approximate agreement between the emulator posterior mean and the true log-likelihood, because those points for

which agreement is poor will be downweighted during importance sampling. Further iterations could instead be performed to improve agreement, but here that was found to be less efficient than having some negligible importance weights. We therefore deem the emulator to be adequate for providing a proposal distribution for the importance sampler.

3.7. *Calibrated output.* We use the emulator estimated at iteration 9 for the final sample of calibrated inputs. We choose this sample to be of size 1,000, and obtain it from an importance sample of size 2,000 by sampling with replacement according to the importance weights, ie. $\exp\{f(x) - m^*(x)\}$. Figure 5 shows the calibrated NHM output against the target data for the four different data types. We can see the calibration to have worked well, as the calibrated output is consistent with the target data, once we account for uncertainty amounts.

3.8. *Sensitivity to the discrepancy specification.* We have calibrated the NHM using discrepancy values of $\lambda_1 = 0.8$, $\lambda_2 = 0.008$, $\lambda_3 = 0.04$ and $\lambda_4 = 0.0004$. We can investigate sensitivity to these choices by simply recalculating log-likelihoods and then importance weights for alternative discrepancy values. This requires little computational cost compared to re-running the simulator. The calibrated output for four alternative discrepancy specifications is shown in Figure 6.

In the first of these alternative discrepancy scenarios, we consider the case where no discrepancy is assumed, which would imply that the simulator is a perfect representation of reality at X^*. This results in an unsatisfactory calibration: all but two of the simulator runs have negligible importance weights, one of which is much larger than the other, and the output from neither of these runs matches the target data, given uncertainty amounts. We then consider doubling discrepancy amounts, relative to our preferred amounts, so that $\lambda_1 = 0.4$, $\lambda_2 = 0.004$, $\lambda_3 = 0.02$ and $\lambda_4 = 0.0002$. This results in the importance sample having a greater range, when compared to the original calibrated inputs of section 3.7, and in turn gives more variability in the calibrated output. While altering the discrepancy specification has changed the distribution of the calibrated inputs, the change in distribution of corresponding output is relatively small, so we do not need to be overly precise when specifying the discrepancy in order to achieve a reliable calibration.

We also consider assuming no discrepancy for only one data source, leaving discrepancy values for the remaining sources unchanged. If we assume no discrepancy for the cases by age data, then the calibrated output still matches the target data for the cases by age data and for the other data

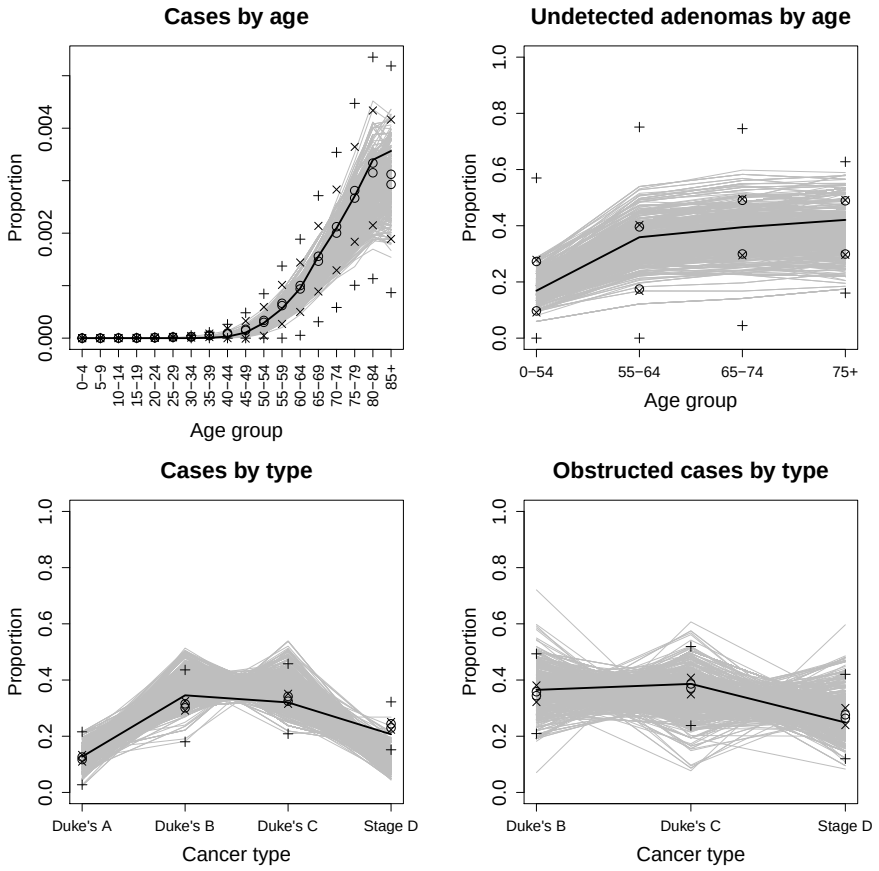

FIG 5. *Calibrated NHM runs against target data.*

sources, and the sample of calibrated inputs also contains sufficiently many unique values. However, when we assume no discrepancy for the cases by type data, the sample of calibrated inputs returns to containing only two unique members (the same two as when no discrepancy is assumed for all data sources), and for cases by type the calibrated output fails to match the target data. In summary we find that, while discrepancy amounts need some consideration, the precision that specifications require is within our capabilities, allowing the NHM to be calibrated reliably. However, the calibration becomes unsatisfactory when we ignore discrepancy, or specify it poorly.

4. Discussion. In this paper we have calibrated a Natural History Model so that its output is consistent with reality. However, in order to

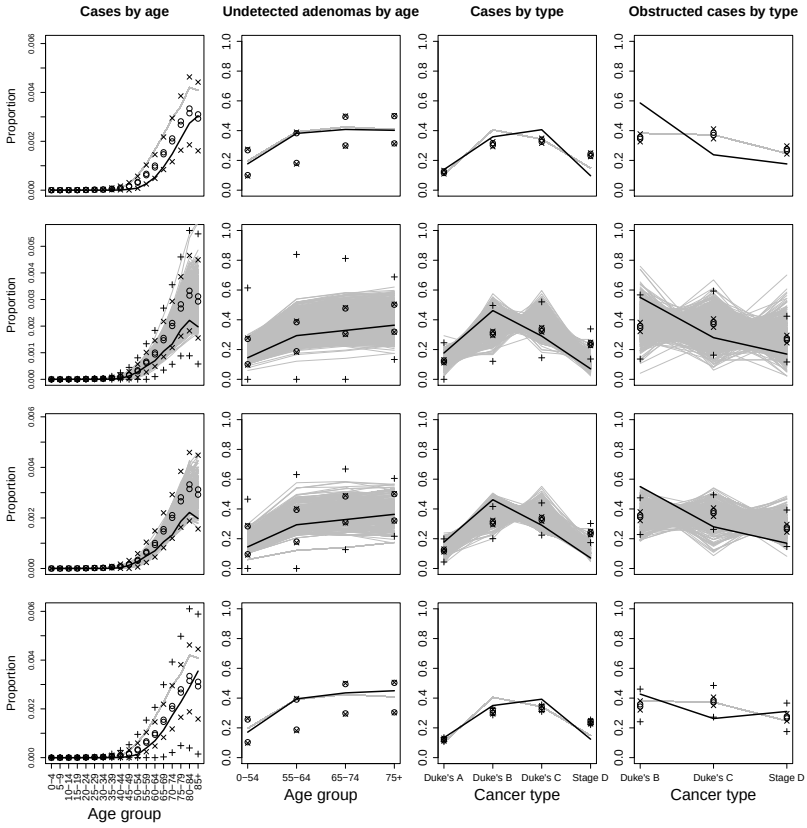

FIG 6. *Summaries of simulator output against target data for various discrepancy specifications: no discrepancy for any data source (row 1), discrepancy levels doubled (row 2), no discrepancy for cases by age (row 3) and no discrepancy for cases by type (row 4).*

do this we have had to address three important issues that arise when calibrating the computer simulator. The first is calibrating a simulator of 'intermediate' computational expense, that is one for which calibration it is not practical using Monte Carlo simulation alone, but nor is it one that requires us to rely solely on a computationally cheap surrogate model, such as a Gaussian process emulator. We therefore propose a calibration method that may be thought of as a hybrid of the two, which uses an emulator to provide a preliminary, approximate calibration, and combines this with simulator run data, through importance sampling, to give a final and more accurate calibration. Because the simulator is only of intermediate computational expense, we have taken a conservative approach to calibration, which can be

seen in the criteria for refining the design region (section 3.4) and when we 'flatten' the log-likelihood (section 3.5). Were the simulator more expensive, we might need to consider optimising the calibration process to minimise the number of simulator runs needed.

The use of importance sampling has allowed us to explore a further issue, which is the sensitivity of calibration to different discrepancy specifications, which is important to understand because discrepancy must be adequately quantified before we can calibrate a simulator [4]. In particular, while we can in theory always adjust a discrepancy specification and check the sensitivity of a calibration to adjustment, in practice this is likely to be impractical due to its computational requirements. Here, though, such investigation becomes computationally feasible, as we simply need to recalculate importance weights and obtain a new sample of calibrated inputs in order to assess different discrepancy specifications. This does need the original importance sample to be suitable, in particular for it to have enough non-negligible importance weights under the new discrepancy specification.

Finally we address how to calibrate a simulator, which we already know to be of intermediate computational expense, that is stochastic and has output that contains count data. We achieve this by using a Gaussian process prior for the log-likelihood, as the log-likelihood is better suited to the Gaussian process assumptions than the simulator output itself. It also reduces the task of calibrating 30-dimensional output to one in which we only have to model a one-dimensional entity. Introducing a nugget effect, overcomes the simulator being stochastic, which will reflect in the log-likelihood surface.

Acknowledgements. We thank Paul Tappenden for proving the NHM, and for guidance on its usage. This work was supported by RCUK funding for the *MUCM2* project (grant EP/H007377/1).

References.

[1] Andrianakis, I. and P. G. Challenor (2012). The effect of the nugget on gaussian process emulators of computer models. *Computational Statistics & Data Analysis 56*(12), 4215–4228.

[2] Bayarri, M. J., J. O. Berger, J. Cafeo, G. Garcia-Donato, F. Liu, J. Palomo, R. J. Parthasarathy, R. Paulo, J. Sacks, and D. Walsh (2007). Computer model validation with functional output. *The Annals of Statistics 35*(5), 1874–1906.

[3] Bayarri, M. J., J. O. Berger, R. Paulo, J. Sacks, J. A. Cafeo, J. Cavendish, C.-H. Lin, and J. Tu (2007). A framework for validation of computer models. *Technometrics 49*(2), 138–154.

[4] Brynjarsdottir, J. and A. O'Hagan (2013). Learning about physical parameters: The importance of model discrepancy. Submitted.

[5] Craig, P. S., M. Goldstein, J. C. Rougier, and A. H. Seheult (2001). Bayesian forecasting for complex systems using computer simulators. *Journal of the American Statistical Association 96*(454), 717–729.

[6] Goldstein, M. and J. Rougier (2006). Bayes linear calibrated prediction for complex systems. *Journal of the American Statistical Association 101*(475), 1132–1143.

[7] Gramacy, R. and H. Lee (2012). Cases for the nugget in modeling computer experiments. *Statistics and Computing 22*, 713–722. 10.1007/s11222-010-9224-x.

[8] Henderson, D. A., R. J. Boys, K. J. Krishnan, C. Lawless, and D. J. Wilkinson (2009). Bayesian emulation and calibration of a stochastic computer model of mitochondrial dna deletions in substantia nigra neurons. *Journal of the American Statistical Association 104*(485), 76–87.

[9] Higdon, D., J. Gattiker, B. Williams, and M. Rightley (2008). Computer model calibration using high-dimensional output. *Journal of the American Statistical Association 103*(482), 570–583.

[10] Higdon, D., M. Kennedy, J. C. Cavendish, J. A. Cafeo, and R. D. Ryne (2004). Combining field data and computer simulations for calibration and prediction. *SIAM J. Sci. Comput. 26*(2), 448–466.

[11] Kennedy, M. C. and A. O'Hagan (2001). Bayesian calibration of computer models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology) 63*(3), 425–464.

[12] Sacks, J., W. J. Welch, T. J. Mitchell, and H. P. Wynn (1989). Design and Analysis of Computer Experiments. *Statistical Science 4*(4), 409–423.

[13] Strong, M., J. E. Oakley, and J. Chilcott (2012). Managing structural uncertainty in health economic decision models: a discrepancy approach. *Journal of the Royal Statistical Society: Series C (Applied Statistics) 61*(1), 25–45.

[14] Tappenden, P. (2011). *A methodological framework for developing whole disease models to inform resource allocation decisions : an application in colorectal cancer*. PhD Thesis, University of Sheffield.

[15] Tekkis, P. P., R. Kinsman, M. R. Thompson, and J. D. Stamatakis (2004). The association of coloproctology of great britain and ireland study of large bowel obstruction caused by colorectal cancer. *Annals of Surgery 240*(1), 76–81.

[16] Vernon, I., M. Goldstein, and R. G. Bower (2010). Galaxy formation : a bayesian uncertainty analysis. *Bayesian analysis. 05*(04), 619–670. This was an invited discussion paper for Bayesian Analysis.

[17] Vernon, I. R. and M. Goldstein (2010). A bayes linear approach to systems biology. *MUCM Technical Report*.

[18] Williams, A. R., B. A. Balasooriya, and D. W. Day (1982). Polyps and cancer of the large bowel: a necropsy study in liverpool. *Gut 23*(10), 835–842.

School of Mathematics and Statistics
The Hicks Building
Hounsfield Road
Sheffield
S3 7RH
UK
E-mail: j.oakley@sheffield.ac.uk

Exeter Climate Systems
College of Engineering, Mathematics and Physical Sciences
University of Exeter
North Park Road
Exeter
EX4 4QF
UK
E-mail: b.youngman@exeter.ac.uk